

**APPLICATION**

**FOR**

**UNITED STATES LETTERS PATENT**

**TITLE:**           **ACCESSING FILE DATA STORED IN NON-  
VOLATILE RE-PROGRAMMABLE SEMICONDUCTOR  
MEMORIES**

**INVENTORS:**   **EDWARD R. RHOADS and JAMES P. KETRENOS**

Express Mail No.: EL 990 136 043 US

Date: March 18, 2004

ACCESSING FILE DATA STORED IN NON-VOLATILE  
RE-PROGRAMMABLE SEMICONDUCTOR MEMORIES

•  
Background

This invention relates generally to processor-based systems using semiconductor memory as their primary, non-volatile, re-programmable storage medium.

5        There is increasing interest in so called embedded processor-based systems. These systems often operate with reduced functionalities to provide desired performance at relatively low cost. In many cases, these embedded systems may be battery operated. Thus, their capabilities may be  
10       limited to improve battery lifetime.

For a variety of reasons including conserving battery life, reducing cost and providing a compact form factor, processor-based systems may be provided which do not use a hard disk drive as their non-volatile storage medium. In  
15       many processor-based system, a hard disk drive provides a convenient non-volatile storage medium that stores most of the information which the user desires to maintain permanently. This may include among other things, the operating system, application software, files and data, as  
20       examples. The information that is stored in the hard disk drive may be transferred for execution to system memory which conventionally is a volatile memory.

In many systems, hard disk drives provide a very high capacity, relatively quick storage medium. However, hard disk drives take more space and use more power than non-volatile semiconductor memories. In many embedded systems, re-programmable, non-volatile semiconductor memories are used as a primary storage system for processor-based systems. These semiconductor memories store the panoply of information normally stored in hard disk drives including operating systems.

In many cases, the semiconductor memories utilized as primary non-volatile storage media for processor-based systems are flash memories. These flash memories may be re-programmed without user intervention using well known on-board capabilities. These memories are generally accessed using row and column addresses. Thus, the memories are generally monolithic in that the location of files and other data in that memory is generally stored outside the memory.

Thus, there is a continuing need for a way to enable an operating system to store more information on a non-volatile re-programmable semiconductor memory and to access that information efficiently.

#### Brief Description of the Drawings

Figure 1 is a schematic depiction of the software modules utilized in accordance with one embodiment of the present invention;

Figure 2 is a schematic depiction of the division of original uncompressed files into blocks in accordance with one embodiment of the present invention;

5 Figure 3 is a schematic depiction of the allocation within a compressed file system image in accordance with one embodiment of the present invention;

Figure 4 is a block diagram of one system for implementing an embodiment of the present invention; and

10 Figure 5 is a flow chart for software which may be used in accordance with one embodiment of the present invention;

Figure 6 is a flow chart for software that may be used in accordance with one embodiment of the invention; and

15 Figure 7 is a flow chart for software for compressing the blocks of information in accordance with one embodiment of the present invention.

#### Detailed Description

Referring to Figure 1, a client processor-based system may include a software architecture 10 having an operating  
20 system kernel 12 that communicates with a file system driver 14. The file system driver 14 receives raw data from a semiconductor memory 40 and arranges that data in a logical layout. The driver 14 communicates with a buffer cache 16 which buffers the raw data to enable it to be  
25 utilized effectively by the driver 14. The device driver 18 accesses blocks of file data from a non-volatile re-

programmable semiconductor memory 40, such as a flash memory, in accordance with one embodiment of the present invention. The device driver 18, which need not have information about the format of the data on the memory 40,  
5 organizes the data in a format that is compatible with the file system driver 14.

Thus, the memory 40 may store a client operating system 42 and a recovery operating system 44 which may be accessed if the client operating system fails. It may also  
10 store a basic input/output system (BIOS) 46 in accordance with one embodiment of the present invention.

The client operating system 42 may include a cyclic recovery check (CRC) field 22, a field 24 that indicates the number of allocation table entries, a field 26 that  
15 includes the allocation table, a field 28 that includes a loader and a field 30 that includes the operating system kernel. The field 30 also includes the drivers 14 and 18.

The client operating system 42 may also have one or more file system data storage areas 34, 36 and 38. These  
20 areas 34, 36 and 38 include raw compressed data that may be utilized by the operating system kernel 12.

The device driver 18 may access any of the data areas 34, 36 or 38 upon request from a file system driver 14. Thus, information may be accessed in the compressed format  
25 on the semiconductor memory 40 and loaded, in an uncompressed format, into the buffer cache 16 for access by

the device driver 14. The device driver 18 decompresses the compressed data in memory 40 and provides it to the buffer cache 16.

5 A compressed file system image may be created initially by forming a single binary file which contains the original uncompressed file system data in its raw form as indicated at 48 in Figure 2. The file containing the original uncompressed file system is then divided into a number of equally sized data blocks 50a-50c. The block  
10 size is the same for each compressed file system image and is set at build time in one embodiment of the invention.

The uncompressed data is then compressed into variable length blocks 58 of compressed data and concatenated together as indicated in Figure 3. Thus, the uncompressed  
15 blocks 50a-50c are compressed to form compressed blocks 58a-58c of the compressed file system image 20.

Each of the areas 34, 36 and 38 (Figure 1) includes an image having a header section (52-56) and a series of compressed blocks 58 which store the file system data, as  
20 shown in Figure 3 in one embodiment of the invention. The header section of the compressed file system image includes a field 52 with cyclic recovery check information. This field may have a zero offset and a length of two bytes. The cyclic recovery check value is calculated over the  
25 length of a block allocation table. The header also includes a field 54 for the number of block allocation

table entries. It has an offset of two and a length of four bytes. The number of entries in the block allocation table may be stored as an unsigned, long value. The actual block allocation table (BAT) 56 has an offset of six and a length which is equal to the number of bat entries. The BAT 56 describes starting offsets and lengths for each compressed block 58.

The device driver 18 uses the BAT 56 to find, in the semiconductor memory 40, the beginning and ending location of each of the compressed data blocks 58. The device driver 18 operates by decompressing the compressed blocks of data in real time and mapping the decompressed data into the file system as requested by the operating system kernel 12 at run time. The device driver 18 may have no knowledge of the file system stored as the compressed file system image 20.

Thus, in accordance with some embodiments of the present invention, an operating system may have access to compressed file system data stored on a semiconductor memory. Semiconductor memories may be less prone to electrical or mechanical failure than hard disk drives. In some embodiments of the present invention, the file system interfaces on the operating system may be utilized and leveraged by application level programs. Since the data stored in the semiconductor memory is compressed, less

memory may be required, resulting in a less expensive solution.

Any file system can be stored in the semiconductor memory independently of the nature of kernel's file system. Thus, the device driver 18 may be unaware of the file system stored within the semiconductor memory 40. In some embodiments of the present invention, additional files may be accessed by the client system 10 from a remote server (not shown). The client may be a processor-based system such as a desktop computer system, a handheld computer system, a processor-based television system, a set-top box, an appliance, a thin client, a cellular telephone or the like. In some embodiments, the system 10 may not be a network connected system.

A storage device implementing the re-programmable semiconductor memory 40 may be electrically re-programmed. The storage device may also act as the BIOS memory for the client in one embodiment of the invention. While conventionally a BIOS memory is a read only memory (ROM), by using a re-programmable memory 40 the operating system as well as the BIOS may be updated or replaced when corrupted. In other embodiments of the present invention, a conventional BIOS ROM may be used in addition to the memory 40.

A variety of flash memories may implement the memory 40, such as Intel's StrataFlash™ brand memory. One



advantageous memory is a 28F640J5 eight megabyte flash array available from Intel Corporation. This memory includes a plurality of one hundred twenty kilobyte blocks. Each block may be data protected so that it may not be  
5 erased or overwritten. In other words, data protection may be selectively applied to one or more of the plurality of blocks in the memory.

A variety of operating systems may be utilized for the kernel 12 including Linux, Microsoft Windows® 98, Windows  
10 2000 or CE, and Be operating systems, as examples. The primary operating system may also be a real time operating system (RTOS) such as the Palm OS® Software 3.5 available from 3Com Corporation.

The recovery operating system 44 operates in cases  
15 where the primary operating system 42 is corrupted or needs updating. The recovery operating system 44 may be an operating system of reduced size which includes basic, essential functions and the limited software needed to obtain a new primary operating system. Thus, as used  
20 herein, a recovery operating system is an operating system that is responsible for updating and/or obtaining a replacement for a primary operating system.

Ideally, the recovery operating system 44 may be stripped down as much as possible to conserve memory. If  
25 possible, its kernel may be reduced to only that code which is necessary to implement its recovery and update

functions. One kernel that is particularly applicable is a Linux kernel. The Linux kernel includes an X-based kernel utility called MakeXConfig. This utility provides a graphical user interface to facilitate selecting the elements of the kernel and the operating system. That is, the Linux operating system allows the user to answer a series of questions, posed through a graphical user interface, indicating whether particular functionalities are desired.

In the case of some system errors or crashes, the client system may reboot thereby resolving the error. If the number of reboots exceeds some threshold level, the recovery operating system may be invoked. When the system attempts to reboot, it may check a CMOS memory reboot count flag and then automatically reboot the recovery operating system if the reboot count threshold is exceeded. The recovery operating system 44 is started so that a new version of the primary operating system 42 may be fetched.

The allocation table (AT) 26 partitions the memory 40 and allows multiple code and data changes to be stored in the memory 40. This in turn allows multiple boot loaders to exist within the memory for booting different operating system images. At boot time, the BIOS 46 may select which boot loader to load and execute based on the status of a recovery bit.

A boot loader 28 for loading the primary operating system may be stored above the allocation table 26. Above the boot loader 28 is the kernel 30 or the core of the primary operating system 42. The primary operating system  
5 42 may be the same or different than the recovery operating system 44.

Above the kernel 30 is the file system. The allocation table 26 includes one entry for each item stored in the memory 40 including the items stored in the file  
10 system. The file system includes files, directories and information used to locate and access operating system files and directories.

Each item contained in the allocation table 26 includes information about the software version, the flags,  
15 the data offsets, the length of the data and its load address. The version number may keep track of which version of software was loaded in a particular memory. The data offset determines where, in the memory 40, an entry is located. The flag field has information about the nature  
20 of the respective entries. The least significant bit of the flag field may include information about the status of the cyclic recovery check. This in effect tells the BIOS whether a CRC must be calculated. The next most significant bit includes the block type. The block type  
25 includes "boot" which indicates a boot loader, "kernel" or "file system". If the block type is boot loader, the flag

field tells where, in random access memory, to load the boot loader out of the memory 40. A boot loader or boot strap loader loads and passes control to another loader program which loads an operating system.

5        While the present invention may be used with a variety of processor-based systems, an application which uses a set-top box as the client system 60 is illustrated in Figure 4. A set-top box works with a television receiver. The client 60 may include a processor 65 coupled to an  
10    accelerated graphics port (AGP) chipset 66. The Accelerated Graphic Port Specification, Rev. 2.0 is available from Intel Corporation, Santa Clara, California. The chipset 66 may be coupled to system memory 68 in the accelerated graphics port bus 70. The bus 70 in turn may  
15    be coupled to a graphics accelerator 72 also coupled to a video or television receiver 73.

      A portion 75 of system memory, called the CMOS memory, may be implemented by memory in integrated circuit which is adapted to save system data. Conventionally, the CMOS  
20    includes a real time clock which keeps the time. Recovery and update bits are stored in the CMOS memory in predefined locations.

      The chipset 66 may also be coupled to a bus 74 and receives a television tuner/capture card 76. The card 76  
25    may be coupled to a television antenna 78 which also may be a satellite or cable connection as additional examples. An

interface to a network 16 such as a modem interface connection to the Internet or a network interface controller to a computer network may also be provided.

5 A bridge 80 may in turn be coupled to another bus 84 which supports a serial input/output interface 86 and a memory interface 94. The interface 86 may be coupled to a modem 88 or a keyboard 92. The interface 94 may couple the memory 40 storing the recovery operating system, BIOS, and primary operating system. The bridge 80 may be the  
10 82371ABPCI ISA IDE Xcelerator (PIIX4) chipset available from Intel Corporation. Thus, it may include a general purpose input/output pins (GP[I,O]).

With a number of chipsets used to implement computer systems, the chipset may be set up so that it only sees a  
15 certain number of lines of BIOS code at any one time. In embodiments in which the primary operating system and the recovery operating system are stored in flash memory, they may be accessed in the same way as BIOS memory is accessed. Thus, since the flash memory that is accessed is  
20 considerably larger than a BIOS memory, it may be desirable to allow other techniques to access all the data stored in the flash memory. One technique for doing this in processors from Intel Corporation is to use the GP[I,O] pins, for example, on the PIIX4 device. These pins can be  
25 coupled to the pins responsible for developing the signals for reading the BIOS. When providing the appropriate

GP[I,O] signals, flash memory reading may be bank switched to sequentially read the entire memory.

The system 60 may create the compressed file system image for storage on the memory. For example, the system  
5 may boot from another storage device, create the file system image and store that image on the memory 40. Alternatively, the image may be created and stored on the memory 40 by an external processor-based system.

Referring now to Figure 5, in accordance with one  
10 embodiment, software 108 begins on power up or system reset with the BIOS executing and performing system initialization and power on self-test activities (block 110). The contents of the memory 40 may be validated by checking the CRC stored in field 96 in the flash memory, as  
15 indicated in block 112. At this point, the BIOS selects the boot loader (block 114) to execute by scanning the allocation table and selecting an entry marked as a boot loader. The boot loader then uses the allocation table to find where in the flash memory the primary operating system  
20 is located (block 116), loads the operating system at the appropriate address in system memory (block 118) and starts its execution (block 120).

Referring to Figure 6, the device driver 18 may begin receiving a request for blocks of data as indicated in  
25 block 124. The requested blocks are accessed from the storage as indicated in block 126. Each block is

decompressed and indicated in block 128. The data is then returned to the file system module as indicated in block 130.

5 The software 132 for compressing the file system image, shown in Figure 7, begins by dividing the file system image into blocks 50 of equal size as indicated at 134. The data is compressed and formed into blocks 50 that are of variable length and concatenated as indicated at 136. The number of entries is determined (block 138) as  
10 well as the CRC (block 140) and the BAT (block 142).

While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended  
15 claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

What is claimed is: